

The “Modeling Clay” Approach to Bio-inspired Electronic Hardware

Ken Hayworth

Jet Propulsion Laboratory, Pasadena CA 91109, USA, Mail-stop 303-300
Ken@brain.jpl.nasa.gov

Abstract. The field of evolvable hardware or bio-inspired hardware holds the promise of automatically engineering complex electronic systems that remain adaptive and fault-tolerant during use. A growing number of experiments along these lines have been performed recently, mostly using off-the-shelf hardware [1] or straightforward extensions of building blocks used by human engineers [2][3]. In this paper we use the POE (Phylogeny Ontogeny Epigenesis) framework of bio-inspired hardware systems [4] and restrict evolutionary search and development considerations to pure hill-climbing search only, in order to develop some theory around evolution of electronic circuits. From this theory a new analog re-configurable hardware architecture is proposed for use in evolvable hardware. The hardware is a context switchable analog computer which can implement any general non-linear dynamic system on the level of the vector field representation. The optimization algorithm is a bio-inspired “molding” of the state-space description of the system. We call this novel hardware/ optimization algorithm platform the “modeling clay” approach to bio-inspired electronic hardware.

1 Introduction

The application of biological and evolutionary concepts to the domain of analog circuits holds the promise of automated analog design as well as promising physical circuits that are more adaptive and robust both in the context of manufacturing errors and changing runtime conditions or damage. This paper asks the question: “What general, re-configurable structure should a piece of evolvable analog hardware have that will allow it to evolve and adapt efficiently?” We seek the answer in biology.

The revolutionary point of Darwin’s theory was quite simple: the slow accumulation of *small random improvements* could, in fact did, produce all the wonderful complexity of the biological world today. To a first, crude approximation, this evolution of complexity can be viewed as a type of hill-climbing search, greedily pursuing all changes that increase an individual’s reproductive success. There is supporting evidence from biological simulations of evolution [7] that literal hill-climbing search can give rise to startling complexity, and this will be reviewed shortly. We do not wish to downplay the clear importance of crossover and the other complexities of population based evolutionary search techniques, we put these important questions aside so as to take advantage of the much simpler viewpoint

which hill-climbing search to complexity allows. This focus on hill-climbing, when looked at within the POE framework of bio-inspired hardware, will direct us toward a new evolvable hardware architecture.

1.1 The POE framework

Sipper et al. [4] introduced the POE framework in which to compare evolvable hardware research directions. POE stands for Phylogeny Ontogeny Epigenesis representing natural selection, embryology, and learning respectively. For this paper we will need to somewhat expand on this framework, using the familiar language of evolutionary theory: "Genome space" is the space representing all possible genomes, and "phenome space", the space representing all possible creatures. If we, for the moment, **restrict** both evolutionary adaptive search and learning to use only greedy hill-climbing search techniques, the POE framework can be viewed as follows:

Phylogeny is the search over genome space using only information from the fitness measure of the corresponding point in phenome space. Evolution represents a hill-climbing search through this genome space. Genome space has a particular topology (a mathematical concept which simply allows us to talk about "connected" paths through these abstract spaces) which for an asexually reproducing creature, may be considered with respect to Hamming distance. Any ancestral tree will be seen to have followed a connected path with respect to the topology of this space.

Ontogeny is the explicit mapping of points in genome space to points in phenome space. Phenome space can be thought of as the space of all possible behaviors. A topology and dimensionality can be imposed on this space as well. Looked at this way, the genome-to-phenome mapping (ontogeny) is a mapping of the entire genome space into a subspace manifold of the phenome space. The key point of this realization is that now the hill-climbing search in genome space can also be viewed as a type of search through phenome space. This is an important point since it is the phenome's fitness that is directing the search in genome space.

We can think of the phenome space as being a road map and the search algorithm as driving around in town. If the genome-to-phenome mapping is "topology preserving" (i.e. points close together in one space should map to points close together in the other, points far from each other should still be far apart preserving connected paths, and the dimensionality of the manifold mapped into should be same as original) then driving to your destination is easy. The less topology preserving the genome to phenome mapping is, the more the road map will lead you astray. (This observation is usually talked about in terms of the "fitness landscape" of the genome space, an equivalent notion but less suited for the following discussion. A smooth fitness landscape is in most ways equivalent to a topology preserving genome-to-phenome mapping, however a fitness landscape based model is always tied to a particular problem to be solved. Since we are endeavoring to design a general piece of analog hardware that will be well suited to a variety of tasks, we must take this more general view. We will expand on this later in the paper and point out some precautions that must be followed when applying it.)

Epigenesis can be thought of as a hill-climbing search over the parameters of the individual phenotype during its lifetime. For instance this might be neural

connections which are determined by the genes but whose values are not. The analysis of this type of learning will require thinking in terms of yet another space, which we will call the "connection space", pushing the neural analogy, of that particular individual phenotype. Simple reinforcement learning is a hill-climbing search through this space, once again using the phenome space's fitness value as a road map. (Reinforcement mediated neural learning algorithms are obviously of this type, gradient decent on energy or error landscapes is the proper analogy.)

Recall that the individual is "born" in the phenome space, at a point which is embedded in the sub-space manifold mapped to by the genome. Epigenesis allows the phenotype to move away from this manifold by changing its "connections" in response to environmental cues. Once again, the success of this search through connection space will depend on how well the connection-space-to-phenome-space mapping preserves topology. A summary is in order:

- Ontogeny is a mapping of genome space into a sub-space manifold of phenome space
- The genome-to-phenome mapping should be topology preserving to allow efficient evolutionary hill-climbing search
- Epigenesis is mediated by a similar hill-climbing search, this time in an individual's "connection space"
- The connection-to-phenome mapping should be topology preserving to allow for efficient adaptation of the individual during its lifetime
- During epigenesis, the individual's point in phenome space, which starts at birth embedded in the sub-space manifold of the genome mapping, is allowed to wander from this manifold into the much larger phenome space accessible through the learning search

The exact definitions of the above spaces and the specifics of the mappings are problem specific, and the above "fuzzy" framework is intended mostly to strengthen intuition. Before we move on to applying this framework to the construction of an evolvable hardware system, let us look at how biology itself fits into this framework, as well as motivate our seemingly severe restriction to use only greedy hill-climbing search.

1.2 An example of an efficient evolutionary system: Mammals

The genome to phenome mapping is the cell differentiation, cell migration, and cell population growth of biological embryology.[6] The genome encodes for proteins. These proteins provide the building blocks for cells, but more importantly many serve as regulatory proteins which, by their concentration, determine cell growth rates and the like. Proteins also regulate other proteins via gene expression, therefore, for each effect in cell behavior leading to embryological development, there is a good possibility that many dozens of genes affect this growth parameter.

Because of the use of regulatory cascades, mutations in specific genes may often produce only quantitative effects like bigger bones or more flattened noses. The reason is that knocking out or altering the shape of one of the proteins in a regulatory network may only affect the *concentration* of the end product, not its form.

One may object that qualitative change cannot occur with only quantitative step changes of concentrations of growth mediating chemicals and cell population sizes. A wonderful simulation (Fig.1) of the evolution of a vertebrate eye [7] brilliantly sets this objection to rest. Also the skeletal structures of some selected mammals [6] showing quantitative variation on a common structure leading to quite different end functions is shown in the fig. 2. It is experiments like that shown in fig. 1 which support the hill-climbing restriction which we are imposing on this discussion.




Fig. 1. Nilsson and Pelger's experiment to prove that cell population size changes alone can produce the structure of a vertebrate eye starting with a simple flat layer of photosensitive cells. (1) proto-eye made up of a flat layer of transparent tissue, covering a flat layer of photosensitive cells, resting on a substrate of black tissue. (2 to 8) progression of evolution to form vertebrate eye structure. Nilsson and Pelger assumed that evolution was a hill-climbing process only and mutations produced small percentage changes in the size of, or refractive index of, small regions of the connected tissue. Beginning eye formed no image, end product "found" the optical solution of the graded index lens. [7][5] **Note how the structure of the problem seems to work well with a hill-climbing technique, this intuition is what we will use in the "modeling clay" approach to electronic evolvable hardware.**



Fig. 2. Drawings of vertebrate limb structure from the illuminating book Triumph of the Embryo [6], reinforcing the idea that simple *quantitative* variations on a common structural scheme can produce **major qualitative** differences

The form of biological embryology has enormous impacts on the efficiency of evolution of biological organisms. It has been said that biological creatures have evolved *to evolve* [8]. Regulatory cascades and smooth tissue-growth variation would form a very important set of constraints on evolution. The embryology of mammalian evolution can be viewed as a set of very strict constraints on the effects of mutations. Symmetry, four limbs, warm-blooded, backbone, etc. are implicit constraints that the

genome to phenome mapping has built in at the mammalian level. When thinking of the incredible variability within the mammals even with these constraints: bats, whales, humans, lions... the “constraints” on the genome-to-phenome mapping can be correctly viewed as a way to make the adaptation via evolution *more* efficient not less. To summarize the lessons learned into the above framework:

- The mammalian genome-to-phenome mapping is highly pre-structured for efficient evolution thanks to billions of years of hard won constraints
- One of these constraints is regulatory cascades of proteins whose end product concentration controls growth rates of tissues. Thus it is possible that many mutations produce only quantitative changes such as bone size and shape differences. This has the effect of making the genome-to-phenome mapping more topology preserving
- Other constraints include symmetry, number of limbs, air breathing lungs, etc, which provide an excellent general body type for a cornucopia of complex creatures: bats, whales, humans, lions...
- The genome space dimension is relatively small compared to the phenome space dimension when the neural epigenesis factor is considered. This allows for evolutionary search for rough body and brain types, and finely detailed specific adaptation during a lifetime.
- Many neural learning algorithms show the mammalian brain uses sensory driven hill-climbing algorithms to efficiently search “connection space”, also showing that the connection-to-phenome mapping may be topology preserving in these cases.

To end this discussion of what we should learn from biological evolution, we will give a quote which helps to reinforce the picture of gradual, continuous change which is characteristic of biological evolution:

“[I]f every form which has ever lived on this earth were suddenly to reappear... it would be quite impossible to give definitions by which each group could be distinguished from other groups, as all would blend together by steps as fine as those between the finest existing varieties...”

-Charles Darwin, as quoted in D. Dennett’s Darwin’s Dangerous Idea p85

2 Bio-inspired electronic circuitry

The term “evolvable hardware” brings to mind electronic components like transistors, resistors, and capacitors, automatically rewiring themselves in response to evolutionary pressure or a reinforcement signal. This vision has driven many researchers to attempt evolvable hardware platforms based on automatic switch-based routing, connecting components [3] or higher level functional blocks [1][2]. It is worth noting that this methodology is an extension of how a human engineer designs circuits. A human engineer starts with discrete components or functional blocks, and has the end goal in mind. This is quite different from what biological evolution does.

The switching in and out of components, in general, does **not** lead to small incremental changes in behavior. Varying component values to effect a smooth transition from one circuit behavior to another often requires pre-determined matched components be tuned together, otherwise the detrimental effects of matching errors will swamp any other improvements (we will show an example of this later). This means that the genome-to-phenome mapping is not topology preserving. This fact forces the researchers to use costly blind-search techniques in an effort to force the genome space topology to higher connectivity (i.e. crossover, multiple mutations, etc.) *compensating for implicit evolutionary cul-de-sacs with an explosion of red herrings*. This requires more of the genome space to be searched than would be necessary if hill-climbing alone could suffice.

We see from this discussion that “bio-inspired” does not simply mean using a GA on an off-the-shelf piece of hardware. It means designing a piece of electronic hardware which has the special properties biological organisms have which makes evolution of complexity possible in the natural world. A few of these special properties were outlined in the previous section. It is time to try to apply these to the field of electronic circuits after we develop one more construct.

2.1 Pseudo-phenome space:

The discussion of the extended POE framework to include hill-climbing search in the various “spaces” may have left the reader a little uneasy, like somehow the “hard part” of the problem was missed. If evolution and learning is simply going toward your goal in phenome space, then what is so hard? Here is the rub; phenome space is the space of all behaviors, an inherently ill-defined concept. Imposing a topology on this space appropriate to, for instance, the six-legged robot walking control problem would require that all methods of locomotion (tripod-gate, caterpillar gate, etc.) lie close to other in this phenome space. (i.e. have connected paths between each other which do not pass through less efficient ways of walking) The prescription from the framework above would then be to “simply” find a genome representation and genome-to-phenome mapping which preserves this topology. A more general problem solving evolvable hardware platform would require a phenome space topology that not only is appropriate for six-legged walking behavior, but also a universe of different behaviors. This is not only difficult but is in principle impossible, and it reaches to the core of how we define “ill-posed” problems.

We believe, however, that the intuition provided by the framework is sound, and that in practice one can design bio-inspired, efficient, evolvable hardware platforms by attempting to approach as close as possible to this ideal. The key to doing this is to define a “pseudo-phenome” space, which is as close to the ideal of behavior topology as possible, and which allows definitions of genome space and connection space mappings preserving the character of this topology. Hill-climbing search is no longer “guaranteed” to reach the optimal solution, but then again, biology has many examples of organisms trapped on local maxima such as whales that have to surface periodically to breathe because of their land-dwelling ancestry. [5] The topology of biological pseudo-phenome space can be imagined by recalling Darwin’s quote above or Nilsson and Pelger’s eye experiment. Tiny changes in tissue growth or synaptic connectivity correspond to small movements in this space. In this paper the




Fig. 3 Active filter circuit from The Art of Electronics p278 [11] along with the three-dimensional state-space of the circuit dynamics with two $Q1 \times Q2$ vector field planes plotted at different points along the x -axis

pseudo-phenome space representation suggested for bio-inspired electronic circuits is the inherently geometric state-space representation of a dynamic system.

2.2 State-space

The *behavior* of an electronic circuit can be described in many ways, but one is universal and precise, the language of differential equations. Any circuit topology can be represented in the following format called state-space form:

$$\begin{aligned}\dot{\vec{q}}(t) &= \vec{f}(\vec{q}(t); \vec{x}(t)) \\ \vec{y}(t) &= \vec{g}(\vec{q}(t); \vec{x}(t))\end{aligned}$$

Where $x(t)$ is a vector of continuous signal values coming into the system, $y(t)$ is a vector of continuous output signal values, and $q(t)$ is a vector of continuous internal state values, the “memory” of the system. The functions $f()$ and $g()$ are vector valued and in general non-linear.

Take, for instance, the circuit in fig 3, which is an active filter design taken out of the book The Art of Electronics [11]. The system is linear and time invariant, has one input and one output. In state-space form the system equations are:

$$\begin{bmatrix} \dot{q}_1 \\ \dot{q}_2 \end{bmatrix} = \begin{bmatrix} 0 & -\frac{1}{R_{F1}C} \\ \frac{1}{R_{F2}C} & -\frac{1}{R_{F2}C} \frac{R_Q}{R_1} \end{bmatrix} \begin{bmatrix} q_1 \\ q_2 \end{bmatrix} + \begin{bmatrix} 0 \\ -\frac{1}{R_{F2}C} \frac{R_Q}{R_G} \end{bmatrix} x \quad (1)$$

$$y = q_2$$

The dynamics of this circuit can be visualized geometrically by constructing a "state-space" which has axes representing the current voltage values of the capacitor states q_1 and q_2 , and the value of the input variable x . As the physical system's currents and voltages change in response to the physics of the circuit connectivity, the point in state-space traces a deterministic curve through the space. At each point in state-space, a vector can be drawn representing dq/dt , this is done by evaluating the matrix equation above at each point in the state space. This produces a fully specified geometric description of the circuit behavior that is in all ways *equivalent* to the differential equation description or the component netlist description.

Try to visualize how the input signal forces the state point through the stack of vector field planes. While the state point is in a particular plane, it is tugged in a particular direction. The rhythmic entrainment of these two forces is what gives rise to the overall circuit behavior, a bandpass filter. We chose this circuit because its filter characteristics (center frequency f_0 , quality factor Q) are qualitative behaviors that can be directly related to all the categories of circuit description (component netlist, differential equations, and state-space vector field). This allows us to compare these various possible representations as possible genome representations.

- **Component netlist:** Most changes in circuit connectivity have enormous impact on the stated behaviors (f_0 , Q), in fact they destroy the filter behavior in one fell swoop. As for changing component *values* in order to smoothly vary the behavior, looking at the circuit you can see that many of the components are pre-matched (those that have the same name). In the original circuit in [11] $R_{F1} = R_{F2} = R_F$, and the circuit allowed both f_0 and Q to be varied independently as can be seen in the equations below when $R_{F1} = R_{F2}$. However without this pre-knowledge we could not vary f_0 without affecting Q , and this could introduce local maxima depending on the fitness measure associated with these behaviors.

$$F_0 = \frac{1}{C\sqrt{R_{F1}R_{F2}}} \quad Q = \sqrt{\frac{R_{F2}}{R_{F1}}} \left(\frac{R_1}{R_Q} \right) \quad (2)$$

- **Differential equations:** If we were to describe the system via the single differential equation below, we would still not be able to set F_0 and Q independently by varying the coefficients A , B , and C .

$$X(t) = A\ddot{Q}(t) + B\dot{Q}(t) + CQ(t), \quad \text{then: } F_0 = \sqrt{\frac{C}{A}}, \quad Q = \frac{\sqrt{CA}}{B} \quad (3)$$

- **State-space vector field:** Using the state-space vector field representation, the higher-level behaviors (f_0 , Q) can be compared to concepts like growth and symmetry which show up in biological embryology. Rewriting (1), and substituting in F_0 and Q for the component values, will allow us to visualize the explicit connections between the behaviors and the vector field geometry.

$$\begin{bmatrix} \dot{q}_1 \\ \dot{q}_2 \end{bmatrix} = \begin{bmatrix} 0 & -F_0 \\ F_0 & -\frac{F_0}{Q} \end{bmatrix} \begin{bmatrix} q_1 \\ q_2 \end{bmatrix} + \begin{bmatrix} 0 \\ -F_0 \frac{G}{Q} \end{bmatrix} x \quad (4)$$




Fig. 4 Relation between circuit behavior and vector field symmetric molding

In this form we can see that increasing the absolute length (growing) of all the vectors in the vector field (along with the vector offset between planes) will increase the center frequency F_0 without effecting the Q . Also, decreasing the rate at which the vectors spiral inward toward the center point (on a particular x value plane) while decreasing the vector offset between planes will increase the Q without effecting the center frequency. This is displayed in figure 4.

Comparing the state-space vector field representation to the other representations, we see that the vector field representation's natural concepts of symmetry allowed the stated behaviors to be de-coupled in the representation. However, the circuit component representation would have required a priori matching of component values, and the differential equation representation would have required breaking the coefficients into more parameters, again a priori. (As an aside, this discussion presents a possible improvement to component representation: a class of mutations that keep resistor pair ratios constant while mutating their value.)

It is important to remember that the vector field representation can be used on non-linear systems as well as linear as in this example. Non-linear stretches of the vector field produce more bizarre transfer functions and quasi-periodic dynamics. Smooth molding of the vector field gives rise to mostly smooth changes in resulting behavior. This is the key ingredient of the state-space vector field representation, which we will use to make an efficient evolvable system.

The space of all state-spaces: Recall that we are interested in creating a pseudo-phenome space in which a natural topology, corresponding roughly to behavior of electronic circuits, exists. The space of all state-spaces (of a particular dimension) meets these requirements. Since the space of all state-spaces is isomorphic to the space of all vector-valued functions, there is a ready distance metric and topology in which to define local search. Following a connected path through the space of all state-spaces should lead you through smooth changes in system behavior, providing more reliable information on which direction to go to reach a desired behavior.

2.3 The “modeling clay” approach defined

With the introduction of the space of all state-spaces as the pseudo-phenome space, we are now ready to apply the extended POE model to bio-inspired electronic circuit theory. Because we now have half a dozen different “spaces” to think about and some of those are “spaces of spaces”, we will dispense with generalizing theories and continue to use the bandpass circuit as a specific example.

Let the genome description be a coarse grained parameterization of common state-space entities. In the bandpass circuit example, the tightness of winding of the vector field, the stretch of the vector field along each of the axes are examples of linear transformations on the vector field. This is not a necessary restriction, but allows us to prove that the resulting genome-to-phenome mapping is a topology preserving mapping from the genome space (a parameterization of the space of linear operators on vector fields) to a sub-space manifold of our pseudo-phenome space: the space of all state-spaces. Since this mapping is a linear transformation from the genome space to the pseudo-phenome space, hill-climbing search for a particular behavior in this parameter genome space will succeed if the search would succeed in the sub-space manifold of the pseudo-phenome space directly. Recall that the pseudo-phenome space was designed so that this is likely the case.

From this phylogeny and ontogeny we get a phenotype embedded in the sub-space manifold of the pseudo-phenome space. In epigenesis, we can use efficient hill-climbing search to travel outward from this useful, but restrictive, “birth” manifold if we define a “connection space” and another topology preserving mapping. For this, break the state space into 1000 blocks, 10 on a side in the three-dimensional bounded state-space. Let the connection space be “control vectors” in each of these boxes, and the connection-to-pseudophenome mapping be a smooth spline interpolant of the resulting vector field defined by the values of the control points.

From spline theory, we see that this is also a linear transformation from connection space onto a set of basis vectors in the pseudo-phenome space. Now reinforcement learning can direct the hill-climbing on the connection space resulting in highly non-linear but smooth stretching and twisting of the dynamic system’s vector field. Thus the final “adult” system can behave much different than the “child” system which was specified by only a few parameters. The full POE framework has been brought into play, and all with efficient hill-climbing search.

We call this the “**modeling clay**” approach because the phylogeny step produces a sculptor’s “rough” through coarse molding and twisting of the vector field, while the epigenesis step details that “rough” with finer, point-directed molding and stretching of the vector field. *The symmetry constraints inherent in the ontogeny mapping, and the spline smoothing enforced in the connection-mapping, make sure that the vector field “behaves like clay” in that local deformations are smoothed out locally, but do not propagate to ruin the hard won beauty which may exist in more distant parts of the sculpture.* Compare this to Nilsson and Pelger’s eye evolution simulation. We have simply found a way of applying the analogy of cell population growth-based *physical* deformation to electronic circuit behavior.

Much of the success or failure of the approach will be dependent on finding molding mutations that, like their embryological tissue growth cousins, produce qualitative changes in phenotypic complexity from only quantitative changes in the pseudo-phenotype structure.

3 Hardware implementation of the “modeling clay”

Besides the promise of hill-climbing searchability toward complex, adaptive behavior, this “modeling clay” approach has a straightforward hardware implementation. An analog computer is an electronic circuit that can embody a set of non-linear differential equations. For instance, at JPL we have recently constructed a simple analog computer that embodies the functionality:

$$\dot{q}_1 = f_1(q_1, q_2), \dot{q}_2 = f_2(q_1, q_2) \quad (5)$$

The functions are bi-cubics. The circuit is constructed such that the coefficients of the terms can be programmed digitally. Thus a wide range of two state variable dynamic systems can be implemented on this computer. What is special about this analog computer, however, is that it is designed to be “context switchable”. Each of the 20 coefficients of the cubic terms is embodied by an MDAC (multiplying D-to-A converter). Each of the 20 MDAC’s are in turn driven by a local digital memory store. This allows the entire analog computer’s dynamics to change via a single broadcast addressing into the multiple MDAC memories. Also, an analog storage stack (not yet implemented) can be included to store and pass the values of the analog state variables. This construction is supposed to mimic software process switching. Fast context switching allows the state-space of a dynamic system to be decomposed into a lookup-table of smaller vector representations. Input, output and internal state values remain analog. Taken together, this is an analog computer architecture that allows for virtual dynamic systems of any size and complexity limited only by memory size. The actual storage of the lookup-table is to be implemented using the CMAC [12] compression and smoothing algorithm. This method uses a combination of overlapping, averaged memories and hash-table storage in order to prevent the memory size from growing exponentially with the dimensionality of the state-space.

The key aspect of this context-switchable analog computer, which is desirable for the “modeling clay” approach to evolvable hardware, is that its *native language* is the vector field representation of dynamic systems. Hardware-in-the-loop phylogeny and epigenesis can be accomplished by allowing the hill-climbing searches to manipulate the memory locations connected to the MDAC’s.

4 Conclusions

Keeping true to Darwin’s vision of the accumulation of small improvements leading to incredible complexity has inspired the design of a novel, bio-inspired electronic hardware. This hardware is a context switchable analog computer which compactly implements an enormous variety of “virtual circuit dynamics” by replacing the “component routing” paradigm dominant in the field, with a new “vector field molding” paradigm. This hardware not only fits within the POE framework, covering all axes of bio-inspired research, but it also allows the POE framework to be seen in a new light where hill-climbing in genome space (P) and connection space (E) is efficient if the space mappings (O) are topology preserving. It is hoped that this extended framework will further intuition in other cases of bio-inspired engineering.

Acknowledgements

The research described in this paper was performed at the Center for Integrated Space Microsystems, Jet Propulsion Laboratory, California Institute of Technology. It was sponsored by the National Aeronautics and Space Administration.

References

1. Thompson, A. et al.: Unconstrained evolution and hard consequences, in Towards Evolvable Hardware Sanchez & Tomassini eds., Springer-Verlag Berlin 1996 p136-165
2. Higuchi, T. et al.: Evolvable hardware and its applications to pattern recognition and fault-tolerant systems, in Towards Evolvable Hardware Sanchez & Tomassini eds., Springer-Verlag Berlin Heidelberg 1996 p118-135
3. Koza, J. et al: Use of architecture-altering operations to dynamically adapt a three-way analog source identification circuit to accommodate a new source, in Genetic Programming 1997 proceedings p213-221
4. Sipper, M. et al.: The POE model of bio-inspired hardware systems, in Genetic Programming 1997 proceedings p510-511
5. Dawkins, R.: Climbing Mount Improbable W. W. Norton & Company NY,NY 1996
6. Wolpert, L. The Triumph of the Embryo Oxford University Press 1991
7. Nilsson, D., Pelger, S.: A pessimistic estimate of the time required for an eye to evolve, Proceedings of the Royal Society of London, B, 256, p53-8
8. Kauffman, S.: The Origins of Order Oxford University Press 1993
9. Dennett, D.: Darwin's Dangerous Idea Simon & Schuster 1995
10. Ritter, H. et al.: Neural Computation and Self Organizing Maps Addison-Wesley 1992
11. Horowitz, P., Winfield, H.: The Art of Electronics 2nd ed Cambridge Univ. Press 1989
12. Albus, J. S.: A new approach to manipulator control: the Cerebellar Model Articulation Controller (CMAC), Trans. of the ASME Sept.1975, p220-7